M.R P. 1051
S.C. 11/234
6th May, 1957

A contribution to the study of Meteorological Time Series

by

J. M. Craddock

(Meteorological Office)

AIR MINISTRY

METEOROLOGICAL RESEARCH COMMITTEE

At the time of issue this report has not been published. Any author desiring to quote this paper should refer to it as "A paper of the Meteorological Research Committee (London),
M.R.P. No. 1051 by Mr. J.M. Craddock.
A copy is available in the Library of the Meteorological Office."

M.R.P. 1051
S.C.II/234
6th May, 1957

AIR MINISTRY

METEOROLOGICAL RESEARCH COMMITTEE

A contribution to the study of Meteorolcgical Time Series

by
J. M. Craddock (Meteorological Office)

Summary

An account is given of the numerical analysis of meteorological time series by means of mutually exclusive band pass filters. It gives methods for producing efficient filters and for the statistical interpretation of the results. Practical examples are given, including an analysis of the slower temperature variations at Kew Observatory, covering the range of periods from 2 days up to about 50 years.

Background

1. This work was undertaken to help to provide a firm statistical foundation to the work of the synoptic meteorologist and the climatologist, so I have given most attention to the features which are important in dealing with the slower atmospheric fluctuations. Nevertheless, the method developed is of far wider application and could be used, for example, in the study of atmospheric turbulence or of tidal series or for the analysis of data regularly distributed in space instead of time.

The outstanding characteristic of the slower atmospheric fluctuations (excluding the strictly periodic ones, based on the day or the year) is their lack of regularity. An oscillation after persisting for several wavelengths may cease abruptly and be followed by another of different period, amplitude and duration. Any method of analysis, such as periodogram or correlogram analysis, which assumes that the series to be analysed preserves its general character throughout a considerable time interval is likely to prove unrewarding when applied to these data, if the statistical properties vary from one part of the series to another. This may happen if the variable is subject to fluctuations with periods as long as or longer than the time covered by the series, or if there are more or less irregular changes in the character of the physical processes producing the observed variations.

The intractability of the material, and the difficulty of assembling and processing the data on the scale necessary to obtain worthwhile results, has meant that very few hard facts have been established about the slower atmospheric fluctuations. Hence the present work has novelty both in the method and in the results. However, the analysis of Kew temperatures though interesting in itself is only one example of many possible applications of a versatile technique.

2. Our problem is parallel to that facing a radio engineer who wishes to design a radio receiver to accept signals from a particular station while excluding signals from other stations, and unorganised noises which may be incident on the aerial at the same time. We can look on our time series as composed of a number of sine waves spread over a particular range of frequencies, but unlike the radio engineer, we do not know initially what part of the range, if any, is likely to

/prove

prove interesting. Our method of analysis consists in splitting the original signal into mutually exclusive components each limited to one section of the original wave band, and measuring the power (or variance) of each component series to find which are the most important. The question whether the power of one component series is significantly greater than that of another is settled by means of Snedecor's F-test.

The actual analysis of the original series is carried out with linear operators, or filters, each of which is designed to pass all oscillations lying in a particular wave band with their original amplitudes and phase relationships, while all oscillations lying outside the band are reduced or annulled altogether. The design of filters is a matter for compromise based on practical experience, since the requirements of efficiency in discrimination and of ability to follow the changes of a non-stationary time series conflict with one another. The filters used here are suitable enough for experimental work, but could be elaborated with advantage when the work is extended.

3. Suppose now that we have a time series $x_1, x_2, x_3 \ldots x_N$ relating to times $\alpha, 2\alpha, 3\alpha \ldots N\alpha$ (N is assumed to be even). Let us carry out a harmonic analysis with basic period $N\alpha$. Then if a_K, b_K are the components of the K^{th} harmonic

$$a_K = \frac{2}{N} \sum_{j=1}^{N} x_j \cos \frac{2\pi j K}{N} \qquad (1)$$

$$b_K = \frac{2}{N} \sum_{j=1}^{N} x_j \sin \frac{2\pi j K}{N} \qquad (2)$$

Any selection of harmonics will give a "least squares" fit to the original data, but if we take all harmonics from $K=1$ to $K = N/2$ then the number of harmonic coefficients will equal the number of data, and the fit will be exact. The frequencies of these harmonics (i.e., the number of cycles per unit time) are $1/N\alpha, 2/N\alpha, 3/N\alpha \ldots \frac{N/2}{N\alpha}$ i.e., they divide the interval $(0, \frac{1}{2\alpha})$ into N/2 equal parts. Thus any time series of N terms at time intervals α can be represented uniquely as the sum of N/2 harmonics equally dividing the interval $(0, \frac{1}{2\alpha})$. This interval does not depend on the number of terms in the series, but only on the time interval between one term and the next, and the higher limit $\frac{1}{2\alpha}$ may be called the cut-off frequency.

Suppose now that the terms of the x series are completely uncorrelated with one another, and that the series is obtained by random drawings from a normal distribution. From equation (1)

$$a_K^2 = \frac{4}{N^2} \left(\sum_{j=1}^{N} x_j \cos \frac{2\pi j K}{N} \right)^2 \qquad (3)$$

and if this is summed over a large number of samples, M

$$\frac{1}{M} \sum^{M} a_K^2 = \frac{4}{N^2} \frac{1}{M} \sum^{M} \left(\sum_{j=1}^{N} x_j \cos \frac{2\pi j K}{N} \right)^2 \qquad (4)$$

Now if the R.H.S. of (4) is expanded, and the order of the summations reversed, expressions such as $\frac{1}{M} \sum x_i x_j$ will vanish in the limit, so that

$$\frac{1}{M} \sum^{M} a_K^2 = \frac{4}{N^2} \sum_{j=1}^{N} \frac{1}{M} \sum^{M} x_j^2 \cos^2 \frac{2\pi j k}{N} \qquad (5)$$

/and

and writing $V(x)$ for the variance of x, etc.

$$V(a_K) = \frac{4}{N^2} V(x) \sum_{j=1}^{N} \cos^2 \frac{2\pi jK}{N} \qquad (6)$$

$$= \frac{2}{N} V(x) \qquad (7)$$

The cos or sin component of the K^{th} harmonic, say $a_K \cos \frac{2\pi Kt}{N\alpha}$ has variance

$$\frac{a_K^2}{2} = \frac{1}{N} V(x) \qquad (8)$$

Thus the average amplitude of any harmonic component obtained by the analysis of a random series of N terms is 1/N times the variance of the original series. Thus while any series of N terms can be represented uniquely by the sum of N harmonic components with equally spaced frequencies $\left(\frac{1}{N\alpha}, \frac{2}{N\alpha} \cdots \frac{1}{2\alpha}\right)$ if the series consists of random terms drawn from a normal distribution the amplitudes of all the harmonics are, on average, equal.

Reverting to our radio illustration it appears that a series which is random, and of which the values of individual terms are therefore wholly unpredictable, appears as a "noise" spread evenly over the whole frequency band. Such a noise cannot be eliminated totally by any method of discrimination between frequencies, whereas it is possible to eliminate unwanted oscillations if these are confined to a band of frequencies which does not overlap the band which interests us. To eliminate an unwanted oscillation we need not reduce its amplitude precisely to zero; it is enough to ensure that the remaining amplitude falls well below the general noise level. This fact reduces the difficulties of filter design mentioned in Section 9.

When a signal is analysed into a series of frequency bands covering the range $\left(0, \frac{1}{2\alpha}\right)$ we must expect it to show significant peaks of power in certain frequency bands rising above a uniform noise level. Conversely, if we have several estimates of the power in different wavebands, all estimates statistically compatible with the lowest estimate of all may be pooled to get an improved estimate of the noise level.

4. All meteorological series contain a random, or noise, component and it is worthwhile considering how this arises. The part which is derived from instrumental errors, or errors involved in rounding observed values to a fixed number of decimal places is truly random, but there is another source which is far more important. To take an example, in Section 15 it is shown that the noise in a long series of Kew daily mean temperatures has variance almost exactly equal to $1°C^2$. In Section 17 the noise variance in Kew monthly mean temperatures is found to be $1.50°C^2$. Of this only about $0.03°C^2$ can originate from the noise of the daily temperatures, averaged in groups of 30, so we have to account for the other $1.47°C^2$.

Now, as shown in Section 3, any series of monthly mean temperatures can be represented as the sum of sine waves of period 2 months (= 60 days) or more, while to represent daily mean temperatures we have to use waves of period 2 days and upwards. Suppose the daily temperatures are represented in this way, and consider the effect of averaging in groups of 30 on the individual sine waves. Waves of period less than 60 days disappear entirely. and are replaced by contributions to longer wavelengths: waves of period somewhat greater than 60 days survive but with reduced amplitudes while waves of period much more than 60 days are hardly affected. The net result is to inflate the noise component in the series of monthly data. In fact, as the figures above show, the noise element in the monthly mean temperatures is nearly all due to the operation of averaging over monthly intervals, which randomises oscillations with frequencies near the cut-off frequency. It follows that if we find that the series we are examining has

/an

an objectionably high noise level, we may be able to reduce the noise level by using the same data but with observations taken at a much shorter time interval and then filtering out the unwanted high frequency variations which before were randomised by the averaging. An example of this is given in Section 24.

5. It may seem artificial to some readers to pay so much attention to the noise component when we know that this is not essentially random at all, but is due to physical processes just as much as the variations we wish to study. However, the noise component, which is unpredictable, and can only be treated statistically, is the price we pay for condensing a multitude of raw data into a far smaller number of averaged data which are easier to handle and to understand. If we overdo the averaging and find that the noise level is objectionably high, we can always go back and use data averaged over shorter time intervals.

The practical design of filters

6. The idea of filtering time series is not new and has been discussed by the Labroustes (1936) who were concerned with filters suitable for use by hand computers, and by Vercelli (1940) who had a mechanical analyser. Unlike these authors, we have electronic computing facilities, and can consider simply what it is worthwhile to do without having to tailor our methods to fit the means of execution. More or less efficient filters have been used by many authors for specific purposes. For example, the ordinary 30-year running mean is a smoothing filter, while Carruthers' periodoscope (1944) is a band pass filter.

7. The operation of filtering, which a radio engineer carries out by means of a tuned circuit, is performed on a time series by means of a moving weighted average. Symmetric filters are additive, in the sense that the sum of the output series from two symmetric filters is exactly equal to the output from the sum of the two filters. This is a great convenience in designing filters, since it enables us to modify an experimental filter progressively towards a desired form, so this paper deals only with symmetric filters.

Let $x_1\ x_2 \cdots x_N$ be our time series and $(w_0, w_1, w_2 \cdots w_K)$ a symmetric filter. Then the operation of generating a new series by the relation

$$y_n = \sum_{j=0}^{K} \frac{1}{2} w_j \left(x_{n-j} + x_{n+j} \right) \qquad (9)$$

is called filtering the x-series with the filter $(w_0, w_1 \cdots w_K)$ and the y-series is the output series. The y-series is shorter than the x-series both at the beginning and the end by K terms. This loss is inherent in the use of symmetric filters, but it may be inconvenient if for example the intention is to predict the extension of the x-series. In this important but rather specialised application, it may be preferable to use asymmetric filters.

8. Consider now the periodic function

$$x = A \cos(2\pi f t + \phi) \qquad (10)$$

for which

$$x_n = A \cos(2\pi f n\alpha + \phi) \qquad (11)$$

Operating on this with the symmetric filter $(w_0, w_1 \cdots w_K)$ gives

$$y_n = \sum_{j=0}^{K} \frac{1}{2} w_j \left\{ A\cos(2\pi f \overline{n-j}\,\alpha + \phi) + A\cos(2\pi f \overline{n+j}\,\alpha + \phi) \right\} \qquad (12)$$

$$= A\cos(2\pi f n\alpha + \phi) \sum_{j=0}^{K} w_j \cos 2\pi f j\alpha \qquad (13)$$

Thus the effect of filtering is to leave unchanged the period and phase of the original oscillation, and to multiply the amplitude by a factor

$$M(f) = \sum_{j=0}^{K} w_j \cos 2\pi f_j \alpha \qquad (14)$$

As shown already, any series with terms at a time interval α can be represented by the sum of harmonics with frequencies not exceeding $\frac{1}{2\alpha}$ so that we need only examine the variations of $M(f)$ for $\left(0 \leq f \leq \frac{1}{2\alpha}\right)$. Further, we can usually choose our time unit to be unity, and then we have to study the variations of

$$M(f) = \sum_{j=0}^{K} w_j \cos 2\pi j f \qquad (15)$$

$$\text{for } \left(0 \leq f \leq \tfrac{1}{2}\right)$$

9. To produce a perfect filter, which passes all waves in a particular frequency band with their true amplitudes and completely eliminates waves of any other frequency we require that M(f) should take specified values at every point of the range $\left(0 \leq f \leq \frac{1}{2}\right)$. Since the symmetric filter $(w_0, w_1, \ldots w_K)$ contains only K + 1 arbitrary constants we cannot satisfy this requirement exactly, but if we choose a reasonable value of K and require M(f) to take the intended value at K + 1 suitably chosen frequencies, we can arrive at a satisfactory approximation.

What I did was to choose a value of K (namely, 6) and then select my K + 1 fitting frequencies spread over the range (0,1/2). Then I considered the K + 1 elementary filters (1, 0, 0 - - - - - 0), (0, 1, 0 - - - - - - 0) .... (0, 0, - - - 0,1) and for each of these wrote down the magnification factor at each of the fitting frequencies. Then I found the K + 1 linear combinations of these elementary filters each of which gives a magnification factor of unity at one of the fitting points and zero at all the others. These combinations (of which Filters No. 71 ----76 and 92 are an example, see Table 2 or 3 or Fig. 9) form a set in which each member passes a rather narrow frequency band and has zero acceptance at the peak frequencies of all the other bands. The output series from any two of the set are almost confined to different frequency bands and are almost uncorrelated, while the sum of the output series from all filters in the set is algebraically equal to the original series. Filters having the property that when applied to a random series they produce uncorrelated output series may be termed orthogonal.

Once we have produced such a set, which represents the maximum resolving power possible with the chosen value of K, the members can be added to produce a filter passing a broader frequency band.

Within limits, the investigator can choose the fitting points to preserve or annul harmonic components known to be present. For example, Filters No. 17 and 117, which annul the 12-term periodicity, are more suitable for use in climatic studies based on monthly mean data than is Filter No. 92 which resembles them except that it annuls the 10-term periodicity.

Finally, if one filter is applied to the output series of another, then the magnification factor for the resultant is equal to the product of those of the two filters. This fact can be used to modify imperfect filters by introducing extra zeros and affords a method of building up complicated filters from simple ones which is less laborious for large values of K than the direct method I have described.

However, I had to remember that the time series I was studying might not be statistically stationary, and decided that the 13-point filters I had produced directly (which are listed in the Appendix) were neither too long nor too short for exploratory work. The frequency response curves for ten typical filters are given in Fig. 9.

/The

The statistical analysis of filtered series

10. Our statistical treatment starts from the null hypothesis that the original series shows nothing but noise and is composed in fact of random drawings from a normal distribution. If then we apply in turn a number of mutually orthogonal filters, each filter will give us an independent estimate of the noise variance, and these estimates may be compared by means of Snedecor's F-test as described, for example, by Brooks and Carruthers (1953).

If $V(x)$, $Vw(y)$ are the variances of the original data assumed to be random and of the M terms of the output series from the filter $(w_0, w_1 \cdots w_K)$ (see equation (9)) then

$$V_w(y) = \frac{1}{M}\sum_{j=1}^{M} y_n^2 \qquad (16)$$

$$= \frac{1}{M}\sum_{j=1}^{M}\left(\tfrac{1}{2}x_{n-K}w_K \cdots + \tfrac{1}{2}x_{n-1}w_1 + x_n w_0 + \tfrac{1}{2}x_{n+1}w_1 \cdots + \tfrac{1}{2}x_{n+K}w_K\right)^2$$

$$= V(x)\left\{w_0^2 + \tfrac{1}{2}\sum_{j=1}^{K} w_j^2\right\} \qquad (17)$$

Thus if we put $P(w) = \left\{w_0^2 + \tfrac{1}{2}\sum_{j=1}^{K} w_j^2\right\}$

and call it the filter factor

then $$V(x) = \frac{V_w(y)}{P(w)} \qquad (18)$$

Thus we obtain an estimate of the variance of the original series on the null hypothesis by measuring the variance of the output series and dividing by the filter factor: this will be called the normalised variance of the filtered series.

11. To compare two such estimates by means of the F-test, we must attach to each the appropriate number of degrees of freedom.

Now a series $x_1 x_2 x_3 \cdots x_N$ derived by random drawings from a normal distribution must have N degrees of freedom. If this series is broken up by a set of orthogonal filters into uncorrelated components which when summed add up to the original series, then the N degrees of freedom must be distributed among the component series,

$$N = N_1 + N_2 \cdots \qquad (19)$$

where the ratio N_1/N is determined only by the weights $(w_0 \cdots w_K)$ of the first filter, the ratio N_2/N by those of the second, and so on.

Now we can look on the N degrees of freedom as being attached not to the terms of the series, but equally well to the N harmonic components by which the series can be uniquely represented: and on the null hypothesis these harmonics are on average of equal amplitude. If now we have a perfect filter, which preserves N_1 of these harmonics while suppressing the remainder, then it seems to me beyond doubt that the number of degrees of freedom of the filtered series is N_1, the number of independent harmonics which enter into it. Our problem, therefore, is to relate N_1 to the weighting coefficients which make up the filter.

/Now

Now for a perfect filter, M(f) is unity for frequencies within the band passed by the filter, and zero elsewhere, and since the harmonics are spaced equally in the frequency range $(0, \frac{1}{2})$

$$\frac{N_1}{N} = \frac{\int_0^{\frac{1}{2}} M(f)\,df}{\int_0^{\frac{1}{2}} 1.\,df} \tag{20}$$

and

$$\frac{N_1}{N} = \frac{\int_0^{\frac{1}{2}} (M(f))^2\,df}{\int_0^{\frac{1}{2}} 1.\,df} \tag{21}$$

Both these expressions give the right answer for a perfect filter, and we have to decide which is right when our filter is not perfect. Bearing in mind that the larger the number of degrees of freedom the larger the expected variance, and that a negative value of M(f), implying a phase reversal at the frequency in question, will still represent a positive contribution to the variance, I think that equation (21) which in these conditions gives a positive contribution to the number of degrees of freedom, must be the correct form.

Thus by equation (15)

$$\frac{N_1}{N} = 2\int_0^{\frac{1}{2}} \left(\sum_{\delta=0}^{K} w_\delta \cos 2\pi_\delta f\right)^2 df \tag{22}$$

$$= P(w) \tag{23}$$

so that $N_1 = N.P(w)$ (24)

This derivation may not seem entirely convincing, but there is no doubt that the problem posed by equation (19) must have a solution, and it is hard to see what the solution can be, other than this.

At any rate, the conclusion leads to sensible results in practice, and the margin of uncertainty, which is present only with imperfect filters, can be reduced by improving the filters. In the present work many of the variance ratios we wish to test range from 5 to over 100, so that a small uncertainty about numbers of degrees of freedom does not matter much.

12. We are now able to split a series up into components which on the null hypothesis are uncorrelated, and to attach the appropriate number of degrees of freedom to each filtered component. This opens the way to the solution of an unsolved problem of great importance in meteorology, namely, that of judging the significance of the correlation between two time series which are themselves autocorrelated. If the original series are analysed by a set of orthogonal filters, and are represented as the algebraic sum of uncorrelated components, then, since it is impossible for an oscillation in one frequency band to be correlated with an oscillation in any other band, the original covariance must be due to a correlation between one or more of the corresponding pairs of components passed by the same filter. It may even happen that the correlation is positive in certain wavebands and negative in others, and that the overall correlation appears to be small although one series is in fact highly dependent on the other.

In any case, the argument of Section 11 enables us to attach a number of degrees of freedom to and hence to judge the significance of the correlation between the output series obtained by the use of the same filter.

/Practical

Practical examples

13. In the next sections, the methods already described will be applied mainly on meteorological data. The actual operation of filtering is laborious to carry out by hand, since a separate calculation consisting of multiplication with summation has to be done to get each term of the filtered series. However, it is very easy to carry out electronically, and all the results given were obtained on the Elliott Bros. 401 Computer operated at the Rothamsted Experimental Station. The machine was programmed to accept a data tape followed by a filter, and to output the filtered series followed by its mean, variance and standard deviation, the estimated variance on the null hypothesis given by equation (20), and the appropriate number of degrees of freedom given by equation (24).

The electronic computation was entirely unchecked, and any machine failure would be liable to affect the results. This risk was taken to save time, since previous work had shown this machine to be extremely reliable. Further, each filtered series has a recognisable character, while the results themselves can be roughly tested by eye. Apart from results affected by known punching errors or vitiated, for example, through the capacity of the machine having been exceeded, none of the output figures gave any suggestion of serious machine failure. I think that there is only the smallest chance that any of the results mentioned here should be misleading owing to an undetected machine fault.

14. The first data consisted of a series of 112 random numbers drawn from an effectively normal distribution. In fact they were made on the Manchester University Mark 1 Electronic Computer, by producing groups of 400 random binary digits from the random number generator, and counting the number of unit digits among them. These numbers should be distributed about a mean of 200 with standard deviation of 10. The details of the analysis are given in Table 1.

Table 1. Filter analysis of random numbers drawn from a normal distribution

Filter no.	Peak frequency	Filter factor	Normalised variance σ^2	Degrees of freedom
17	0	.0798	8.5	8
31	1/12	.1527	10.5	15
21	1/6	.1528	6.3	15
35	1/4	.1528	12.2	15
14	1/3	.1528	10.8	15
33	5/12	.1527	13.8	15
24	1/2	.0798	14.1	8

It will be seen that none of the normalised variances differs significantly from the expected value, 10, nor does the ratio of any one to any other exceed even the 5 per cent significance level. This analysis gives no ground whatever for doubting the null hypothesis that the original series is as random as it was intended to be. The original series and the filtered series are plotted out in Fig. 1 to give an impression of the appearance of series covering these particular wavebands. The series passed by Filter No. 24 shows curious variations in amplitude suggesting a real wave with period nearly equal to twice the interval between the terms. However, all series passed by this filter have this appearance, which must be mathematical in origin and cannot be taken as evidence of the physical existence of such a wave.

15. An analysis was made of daily mean temperature anomalies at Kew. 377 daily values were taken, each the mean of 24 hourly observations, covering the year 1955 with six extra terms at each end, so that the filtered series should cover the calendar year. The annual and semi-annual variations were removed by subtracting daily normals obtained by fitting a two term harmonic form to mean temperatures averaged over the 50 years 1901-1950. This should not affect the frequencies studied in the present analysis, of which details are given in Table 2.

/Table 2.

Table 2. Filter analysis of Kew daily mean temperatures for 1954

Filter no.	Peak frequency	Filter factor	Normalised variance $^{o}C^{2}$	Degrees of freedom
71	$\frac{1}{2}$	.0852	0.95	31
72	$2/5$	.1493	1.03	54
73	$1/3$	.1516	1.64	55
74	$\frac{1}{4}$	.1502	2.49	55
75	$1/6$	.1516	5.39	55
76	$1/10$	.1493	13.63	54
92	0	.0852	46.59	31

The normalised variances strongly contradict the null hypothesis, in that while the outputs for the higher frequencies are mutually compatible, those for all longer periods are not. They show that significant atmospheric fluctuations exist with peak frequency $1/10$ or less (i.e., with period 10 days or more).

The output from Filters No. 71 and 72 can be pooled to give an estimate of the noise variance of Kew daily mean temperatures of $1.00^{o}C^{2}$, based on 85 degrees of freedom. Sections of some of the filtered series are plotted out in Fig. 2.

16. The next stage was to analyse Kew temperatures averaged over non-overlapping periods of 5 days. Again, a series of 377 terms was used, covering the five years 1946-50 with six extra terms at each end. The daily values were estimated as $\frac{\max + \min}{2}$ and the annual and semi-annual variations were removed as before. Details of the analysis are given in Table 3, and sections of the filtered series are given in Fig. 3.

Table 3. Filter analysis of Kew 5-day mean temperatures for 1946-1950

Filter no.	Peak frequency	Filter factor	Normalised variance $^{o}C^{2}$	Degrees of freedom
71	$\frac{1}{2}$	.0852	2.79	31
72	$2/5$	.1493	2.15	54
73	$1/3$	.1516	3.27	55
74	$\frac{1}{4}$	.1502	4.31	55
75	$1/6$	.1516	7.92	55
76	$1/10$	.1493	10.55	54
92	0	.0852	20.48	31

These variances again contradict the null hypothesis and show that there are significant atmospheric fluctuations with periods of 30 to 50 days or more.

17. The analysis is continued on Kew monthly mean temperatures for the years 1921-50 given by Marshall (1952), and is summarised in Table 4.

Table 4. Filter analysis of Kew monthly mean temperatures for 1921-1950

Filter no.	Peak frequency	Filter factor	Normalised variance $^{o}C^{2}$	Degrees of freedom
24	$\frac{1}{2}$	.0798	1.35	28
33	$5/12$	.1527	1.43	53
14	$1/3$	.1528	1.63	53
35	$\frac{1}{4}$	.1528	1.53	53
21	$1/6$	.1528	3.69	53
31	$1/12$	.1527	151.7	53
17	0	.0798	3.21	28

/It

It appears that the variance in the bands centred on periods of 2, 2.4, 3 and 4 months are compatible, and can be pooled, giving a joint estimate of the variance of the random component of Kew monthly mean temperature, based on 187 degrees of freedom of 1.502°C^2.

The 6-monthly component, the semi-annual variation is significantly greater than this (at the 1 per cent level) while the annual variation is much greater still. The components of longer period passed by Filter No. 92, which may be termed the interannual variations, are also significantly larger than expectation, in about the same ratio as the semi-annual variation.

18. Kew monthly temperatures for the 86 years 1871-1956 taken from Marshall (1932) were averaged over 3-monthly periods to give the next series for analysis, of which the results are summarised in Table 5.

Table 5. Filter analysis of Kew 3-monthly mean temperatures for the years 1871-1956

Filter no.	Peak frequencies	Filter factor	Normalised variance °C^2	Degrees of freedom
40	$\frac{1}{2}$, $\frac{1}{4}$	.2465	77.00	81
41	5/12, 1/3	.3334	.75	110
21	1/6	.1528	.91	50
31	1/12	.1527	.99	50

The special Filter No. 40 used on 3-monthly data passes the annual and semi-annual variation with their true amplitudes. Filter No. 41 passes a band of periods between about 7 and 9 months, which is very hard to examine because of its proximity to two much larger oscillations. However, it does not appear to have significantly greater power than the bands centred around 18 months or 36 months and these estimates can be pooled to give one estimate of the variance of 3-monthly mean temperatures at Kew of 0.843°C^2 based on 210 degrees of freedom.

19. Finally, the analysis was applied to the annual mean temperatures for Central England for the years 1698-1952 as estimated by Manley (1953). Manley's data were used instead of the shorter series of annual mean temperatures at Kew, because by making the numbers of degrees of freedom as large as possible the maximum resolution is obtained among the fluctuations on the longer time scales. Details are given in Table 6.

Table 6. Filter analysis of annual mean temperatures estimated by Manley (1953) for Central England for the years 1698-1952

Filter no.	Peak frequency	Filter factor	Normalised variance °C^2	Degrees of freedom
24	$\frac{1}{2}$	.0798	.421	19
33	5/12	.1527	.244	37
14	1/3	.1528	.288	37
35	$\frac{1}{4}$	.1528	.239	37
21	1/6	.1528	.299	37
31	1/12	.1527	.415	37
17	0	.0798	.601	19

The second to fifth variances are obviously compatible, and when pooled these give an estimate of the variance of the random element in the annual mean temperatures of 0.267°C^2 based on 148 degrees of freedom. The output of Filter No. 24 has a ratio to this estimate of 1.58 which is about on the 5 per cent level for significance. The output of Filter No. 31 has a ratio of 1.55 which is significant at a level between 5 and 1 per cent. The output of Filter No. 17 has a ratio of 2.25 which is significant at the 1 per cent level.

These annual data thus provide three features of interest, namely

(1) a tendency for alternation of years. This is only just significant,

and, until supporting evidence is produced, it would be as well to treat it with reserve

(2) significant oscillations are passed by Filter No. 31 which has peak frequency at 1/12, and zeros at 1/6 and 0. If there were significant temperature oscillations related to the 11-year sunspot fluctuation, they would certainly show up in this frequency band, but it would be unjustifiable at this stage to assume the existence of a causal relationship; the question can be tested by finding the correlation between this filtered series and a similar series of filtered sunspot numbers and judging its significance as suggested in Section 12

(3) there is a significant variation on the longest time scale, passed by Filter No. 17, which may reasonably be regarded as climatic change. The non-normalised variance passed by this filter is only $0.048°C^2$, so that the climatic change is not very impressive in magnitude.

20. The variations of Kew temperatures have been discussed for several different wavebands, and I will now try to fit them all together. It is not difficult to show that a 13-point filter with factor P acting on m-day mean data produces the same variance as a 13m-point filter with factor P/m acting on daily data; thus all the estimates of variance can be rendered comparable by multiplying them by the appropriate value of m.

These estimates can generally be plotted against the peak frequency of the filter which relates to them, but it will be seen that since, for example, Filter No. 21 which has peak frequency 1/6 will pass half or more of any oscillation in the frequency range (1/5, 1/10) there is a margin of uncertainty as to what really produces the variance in question. Also, as regards the lowest frequencies of all, the output of Filter No. 17 which passes the oscillations with period much exceeding 12 years must be allotted to some frequency within the range which can be resolved by a series of data covering no more than 255 years. In all cases, uncertainty about the details of the oscillations producing any estimate of variance can be removed by simply plotting the filtered series.

The normalised variances vary in the range of 1 to 250, while the interesting variations are all at the low frequency end of the spectrum. Consequently, a plot of log V_n against log f displays the information more clearly than would a linear plot, and is used in Fig. 4.

Of the 28 points, 25 lie very nearly on a smooth curve. The other three give the familiar annual and semi-annual variation, and the biennial variation mentioned in Section 21. If this biennial variation were spurious, due to the randomising of higher frequency variations in forming the annual means, then the adjacent points on the curve would be affected in the same way, whereas in fact they agree among themselves on a lower value of the variance in this region, so with some hesitation I am inclined to regard the biennial variation as real.

If a series of random numbers was analysed in the same way the resulting graph would be horizontal, so that the presence of horizontal stretches in the graph in Fig. 4 would indicate the existence of frequency bands not characterising the operation of any specific physical processes, i.e., bands in which the variations were due entirely to the chance superposition of oscillations of higher frequency. In fact the graph does not encourage us to think that there are any such flat regions, but suggests that positive contributions to variation are occurring on every possible time scale.

21. While Fig. 4 is convenient for showing variances on every time scale, it does not show them in true relation of magnitude. To do this we plot V_n against log f and obtain an "equal area" graph in the sense that frequency bands enclosing equal areas under the curve will correspond to oscillations of equal amplitude. This is done in Fig. 5. This shows quite clearly that the most important temperature fluctuations at Kew lie in the middle band of frequencies, in the range of periods from about 5 to 60 days. The smooth curve drawn among the points is obtained by transforming that in Fig. 4. Apart from the annual, semi-annual and biennial components mentioned already, the smooth curve gives a fairly good

/fit

fit to the wings of the curve, but near the peak it gives values which differ considerably from those observed. It is difficult to avoid the conclusion that the power spectrum at Kew has two peaks centred about periods of 10 and 35 days. The dotted curve drawn freehand to follow the observations is probably nearer to the truth. It is clear in any event that the peak is a fairly broad one, and that (excluding the astronomical components) perhaps two thirds of the variance of atmospheric temperature fluctuations at Kew is associated with periods of between 5 and 60 days.

The lower peak at about 10 days may depend on the movement of long waves in the westerlies. As regards the peak at 35 days, it appears from work by Craddock and Lowndes (1956) that fluctuations on about this time scale are large in amplitude and in geographical extent over the greater part of the Northern Hemisphere, and that they appear to be associated with abnormalities in the general circulation of the westerlies. There is thus a presumption that at Kew apart from the annual and semi-annual variations the largest fluctuations to which the atmosphere is subject are associated with two types of variation in the general circulation.

22. It is interesting to compare Fig. 5 with a very similar graph published by Griffith, Panofsky and van der Hoven (1956) based on data for University Park, Pennsylvania. Their graph shows peaks at about 12, 38 and 900 days of which the first two probably correspond to our peaks at 10 and 35 days.

The peak at 900 days is not beyond doubt, because the American low frequency values are based on data for a much shorter period of years than ours, and are subject to wider sampling fluctuations. However, it may correspond to our peak at about 730 days.

The American results also give a very pronounced peak at about 4 days of which there is no sign in Fig. 5. This is not due to any fault in the present calculation because, as Fig. 2 shows, the filtered series show that at Kew there is no peak amplitude at this frequency. Also comparing the peaks near 10 and 35 days, the second is relatively larger at Kew than it is in Pennsylvania.

These differences point to deep seated differences between the workings of the atmospheric heat engine at the two places which, in view of the differences in latitude and geographical situation, is not really surprising.

23. To finish this discussion of Kew temperatures Fig. 6 and 7 show the actual course of the fluctuations in the wave bands which are most interesting. Bearing in mind that the fluctuations must lie within certain limits to be passed by the filter at all, they show almost as much irregularity as it is possible to imagine. There is no consistency either in period or amplitude, although on average the amplitudes are so large that it is impossible for any meteorologist to ignore them.

The prediction of the course of such intractable series is bound to be difficult, whether looked at from the physical or the statistical standpoint, but is made somewhat less formidable by the fact that the longer of these oscillations are so slow that a forecast covering no more than a single wavelength would have practical value. This topic, however, is outside the scope of the present paper.

24. A final example shows the value of filtering analysis in the study of climatic change. The basic data are monthly mean temperatures for Bear Island for the years 1923-June, 1941. These observations cover only $18\frac{1}{2}$ years and if annual averages are taken only 18 degrees of freedom will be left, far too few for us to be able to find anything of interest. However, we filter the monthly mean temperatures for the $18\frac{1}{2}$ years with three filters, one passing waves of period up to four months, the next the semi-annual and the annual period and the third only the waves of period exceeding 12 months. In the first and last the annual and semi-annual waves are annulled entirely. The results are given in Table 7, and the series are displayed in Fig. 8.

/Table 7.

Table 7. Filter analysis of Bear Island monthly mean temperatures to show inter-annual variations

Filter no.	Peak frequencies	Filter factor	Normalised variance $^{\circ}C^2$	Degrees of freedom
115	$\frac{1}{2}$, $\frac{2}{5}$, $\frac{1}{3}$, $\frac{1}{4}$	.5879	3.52	123
114	$\frac{1}{6}$, $\frac{1}{12}$	.3334	58.85	70
117	0	.0797	15.60	17

It appears that the variance in the longer wavelengths is decisively larger than the chance expectation, so that we have evidence of significant inter-annual temperature variations. There is a minimum in the late 1920's and a maximum in the late 1930's followed by a sharp fall. There is also a strong suggestion of an oscillation of smaller amplitude of period 3 to 4 years.

These features deserve further study but need not delay us here: the important fact is that it has been possible to identify them at all in data covering so short a period.

Conclusion

25. The main object of this work has been to describe a method of analysis applicable to a very wide range of meteorological time series without assuming the series to be stationary in the statistical sense. The features for which novelty is claimed are

(1) the use of orthogonal filters covering the whole range of possible frequencies,

(2) the comparison of the variance of the output from different filters by means of the variance ratio test, and

(3) the use of comparatively short filters to avoid obscuring the transient features of a non-stationary time series.

The literature shows a dearth of worthwhile analyses of meteorological data and the whole project, both theoretical and practical, had to be planned with very little prior knowledge of the probable outcome. Nevertheless, the results though intended primarily to indicate the flexibility of the method are of considerable practical importance.

To the meteorologist they indicate the time scales on which long range forecasting research is most likely to yield worthwhile returns, while they show the statistician that the mass of statistical theory concerned with stationary time series is not readily applicable to the main problems of synoptic meteorology.

Acknowledgments

Acknowledgments are due to Messrs. M. J. R. Healy and Mr. D. Rees of Rothamsted Experimental Station for assistance in getting the electronic programmes to work and in carrying out the machine calculations.

References

Brooks and Carruthers	1953	Handbook of Statistical Methods to Meteorology, M.O.538 H.M.Stationery Office, London, 1953
Carruthers, N.	1944	Quart. J. R. Met. Soc., London, Vol. 70, p.141
Craddock and Lowndes	1956	Meteorological Research Papers, Air Ministry, London, No. 994
Griffith, Panofsky and van der Hoven	1956	J. Met., Lancaster, Pennsylvania, Vol.13, pp. 279-282
Labrouste, H. and Mrs.	1938	Terr.Mag., London, Vol.XLI, pp.15-28 & 105-126
Manley, G.	1953	Quart. J. R. Met. Soc., London, Vol.79, pp.242-261
Marshall	1952	A Century of London Weather
Vercelli, F.	1940	R. Comitato Talassografico Italiano, Venezia, Memoria 285

Appendix. Sets of 13-point band pass filters used in the analysis

Each of these sets has the following properties

(1) when any two of the set are applied in turn to a series obtained by random drawings from a normal distribution the two output series are almost uncorrelated (i.e., the members of the set are mutually almost orthogonal)

(2) each filter has zero acceptance at the peak frequencies of all other members of the set and unit acceptance at its peak frequencies and

(3) when all members of the set are applied to one data series in turn, the algebraic sum of all the output series is equal to the original series.

Set 1

Filter refce. number	Coefficients							Peak frequency	Filter factor
71	.1063	-.2127	.1897	-.1667	.1437	-.1206	.0603	$\frac{1}{2}$	.0852
72	.1708	-.2764	.1708	0	-.1708	.2764	-.1708	$2/5$	.1493
73	.1273	-.1273	-.2060	.3333	-.1273	-.2060	.2060	$1/3$	.1516
74	.1910	0	-.3090	0	.3090	0	-.1910	$\frac{1}{4}$	.1502
75	.1173	.1273	-.2060	-.3333	-.1273	.2060	.2060	$1/6$	.1516
76	.1708	.2764	.1708	0	-.1708	-.2764	-.1708	$1/10$	.1493
92	.1065	.2127	.1897	.1667	.1435	.1206	.0603	0	.0852

Set 2

Filter refce. number	Coefficients							Peak frequency	Filter factor
24	.0833	-.1667	.1667	-.1667	.1667	-.1667	.0833	$\frac{1}{2}$	.0798
33	.1667	-.2887	.1667		-.1667	.2887	-.1667	$5/12$	.1527
14	.1667	-.1667	-.1667	.3333	-.1667	-.1667	.1667	$1/3$	.1528
35	.1667	0	-.3333	0	.3333	0	-.1667	$\frac{1}{4}$	.1528
21	.1667	.1667	-.1667	-.3333	-.1667	.1667	.1667	$1/6$	.1528
31	.1667	.2887	.1667	0	-.1667	-.2887	-.1667	$1/12$	.1527
17	.0833	.1667	.1667	.1667	.1667	.1667	.0833	0	.0798

Set 3

Filter refce. number	Coefficients							Peak frequency	Filter factor
40	.2500	-.1667	-.1667	-.1667	.5000	-.1667	-.0833	$\frac{1}{4}$, $\frac{1}{2}$	.2465
41	.3333	-.4553	0	.3333	-.3333	.1220	0	$1/3$, $5/12$	.3334

These two filters form an orthogonal set with Nos. 21, 31 and 17 of Set 2.

Set 4

Filter refce. number	Coefficients							Peak frequency	Filter factor
114	.3317	.4581	-.0016	-.3333	-.3317	-.1248	.0016	$1/6$, $1/12$	.3334
115	.5874	-.6293	-.1625	.1666	.1626	-.0373	-.0875	$\frac{1}{2}$, $2/5$, $1/3$, $\frac{1}{4}$	.5879
117	.0809	.1712	.1641	.1667	.1691	.1621	.0859	0	.0797

Notation and Definitions

If $x_1, x_2 \cdots x_n \cdots$ is any time series and $(w_0 w_1 \cdots w_K)$ is a filter, then the series $y_1 y_2 \cdots y_n \cdots$ defined by the equation

$$y_n = \sum_{j=0}^{K} \tfrac{1}{2} w_j (x_{n-j} + x_{n+j}) \quad \cdots\cdots (1)$$

is the output from the series $x_1, x_2 \cdots x_n \cdots$ from the filter $(w_0 w_1 \cdots w_K)$.

If the x-series is strictly periodic, so that

$$x = A \cos(2\pi f t + \phi)$$

then

$$y_n = x_n . M(f)$$

where

$$M(f) = \sum_{j=0}^{K} w_j \cos 2\pi j f \quad \cdots\cdots (4)$$

$M(f)$ is called the magnification factor of the filter $(w_0, w_1 \cdots w_K)$ and $M(f)$ plotted against f for $(0 \le f \le \frac{1}{2})$ is called the response curve of the filter.

If $V(x)$, $V_w(y)$ are the variances of the original data, (assumed to be a series of random numbers), and of the output series from filter $(w_0, w_1 \cdots w_K)$, then

$$V_w(y) = V(x) . P(w)$$

where

$$P(w) = \left\{ w_0^2 + \tfrac{1}{2} \sum_{j=1}^{K} w_j^2 \right\} \quad \cdots\cdots (5)$$

and

$$V(x) = \frac{V_w(y)}{P(w)} \quad \cdots\cdots (6)$$

$P(w)$ is called the filter factor and $V(x)$ estimated in this way is the normalised variance of the output series.

When a data series is operated on in turn by each of several mutually exclusive filters, then on the null hypothesis the various normalised variances are independent estimates of the same thing.

Statistical Problems

If we have N terms of the output series from filter $(\omega_0, \omega_1, \ldots \omega_K)$ and the variance measured from these has N_1 degrees of freedom, how do we express N_1 as a function of N and $(\omega_0 \omega_1, \ldots \omega_K)$?

I suggest that

$$\frac{N_1}{N} = 2\int_0^{\frac{1}{2}} (M(f))^2 df$$

which gives

$$N_1 = N.P(\omega) \quad \cdots\cdots\cdots (10)$$

This seems to be right if we have perfect filters. I assume without proof that it holds also for imperfect filters. We can, however, design filters to approach perfection as nearly as we please by taking enough trouble.

Is any adjustment necessary because to get N terms of filtered series we require $N+2K$ terms of original data?

The null hypothesis is that the original data were obtained by random drawings from a normal distribution. How should we modify the analysis if the frequency distribution of the values making up the time series is not even approximately normal?

Daily rainfall or sunshine totals and daily or yearly totals of sunspot numbers form series of this type.

What features do we expect the power spectrum to show if the original data are drawn at random from a distribution which is not normal?

FIG. 1 A SERIES OF NUMBERS DRAWN AT RANDOM FROM NORMAL DISTRIBUTION ANALYSIS WITH FILTERS NOS. 24, 33, 14, 35, 21, 31 & 17.

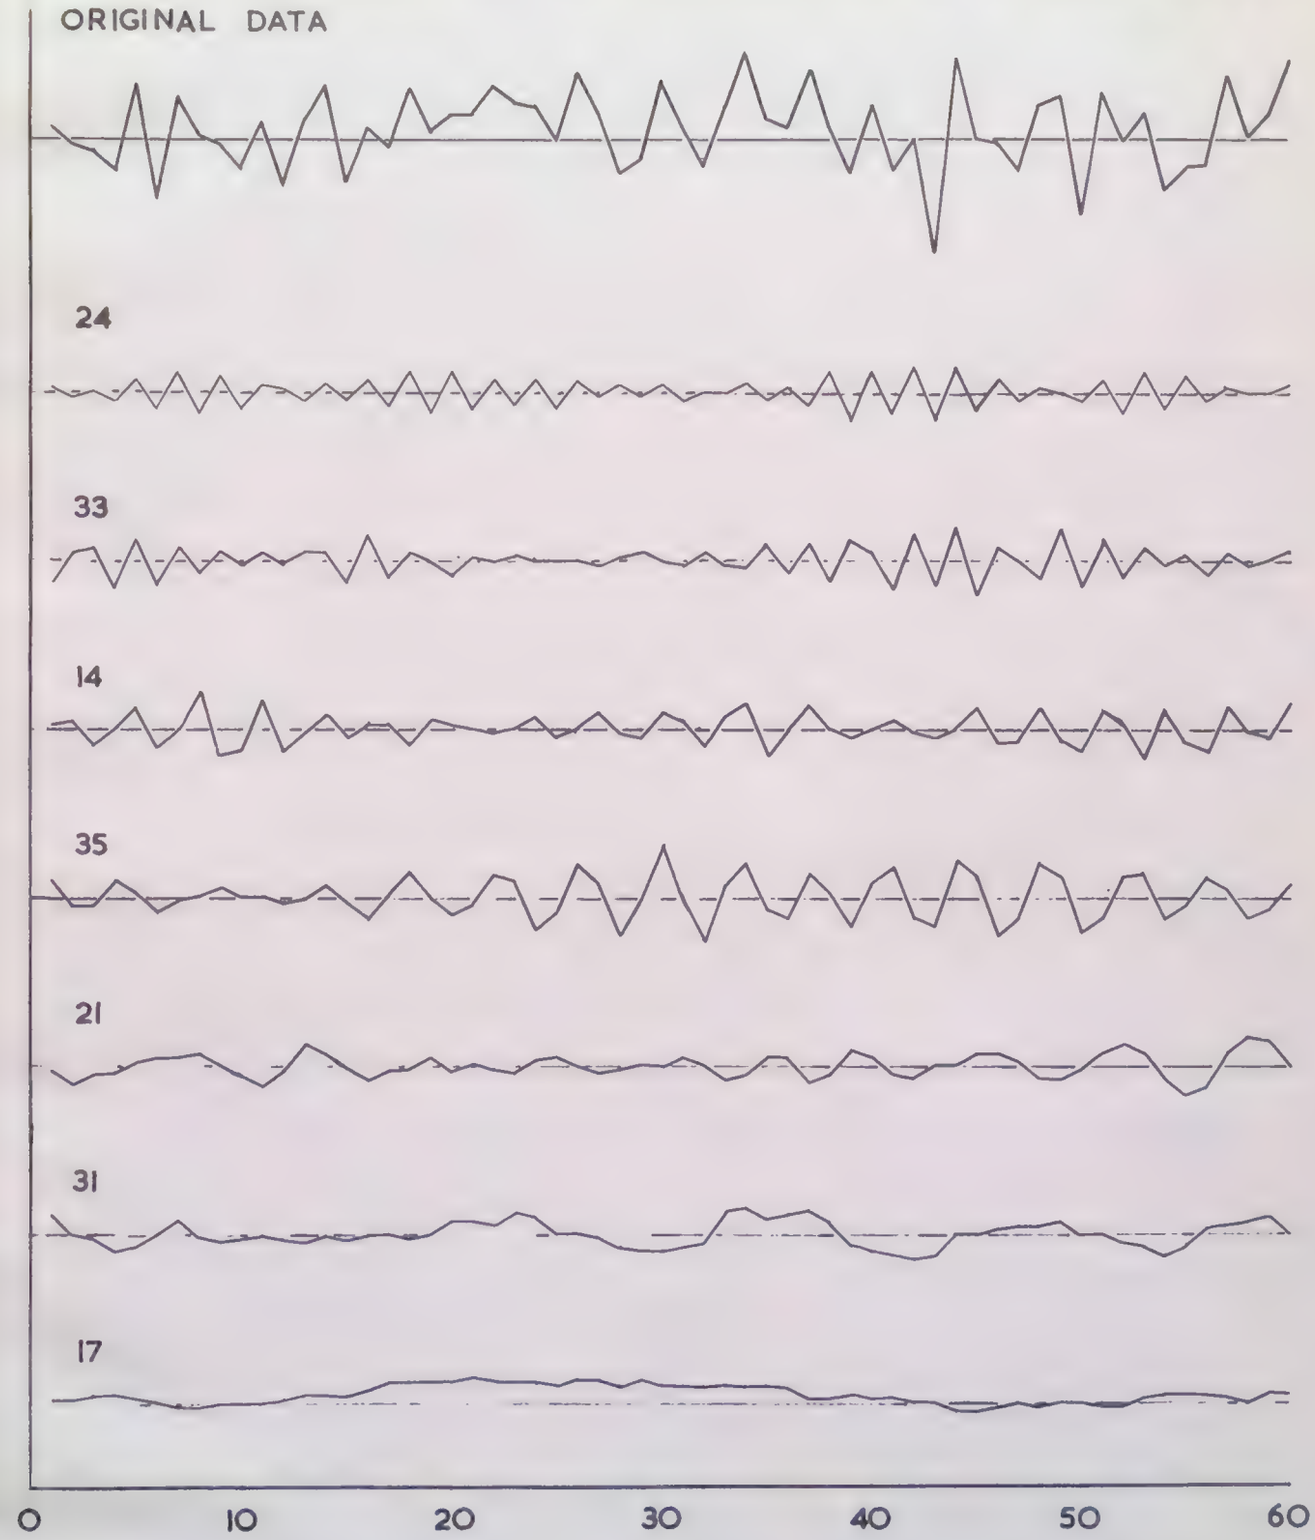

M.R.P. 1051

FIG. 2. FILTER ANALYSIS OF KEW DAILY MEAN TEMPERATURES
FOR JANUARY AND FEBRUARY 1955 .
REFERENCE NUMBERS OF FILTERS GIVEN IN BRACKETS.

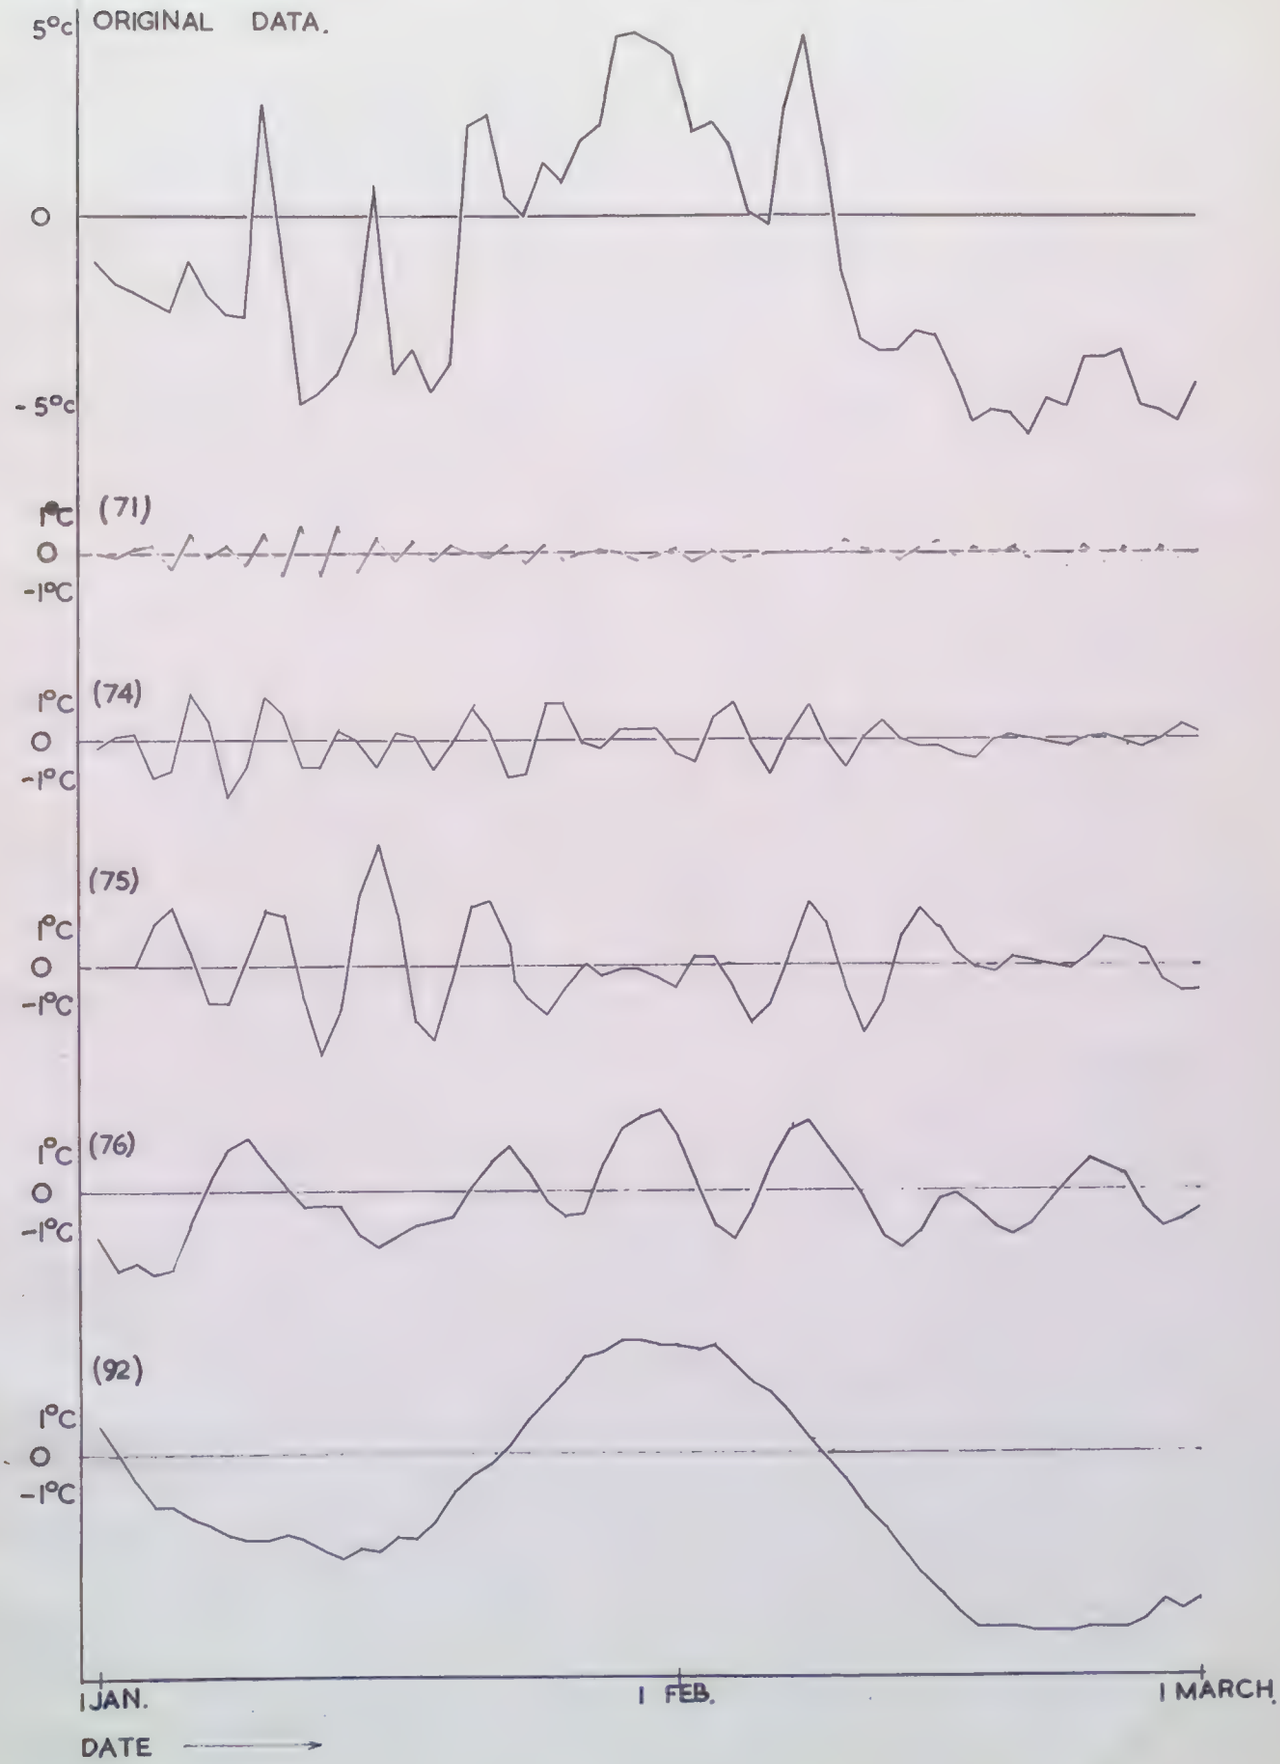

Fig 3. Filter Analysis of Kew 5-day mean temperatures for Jan - Oct 1946
Filter numbers given in margin.

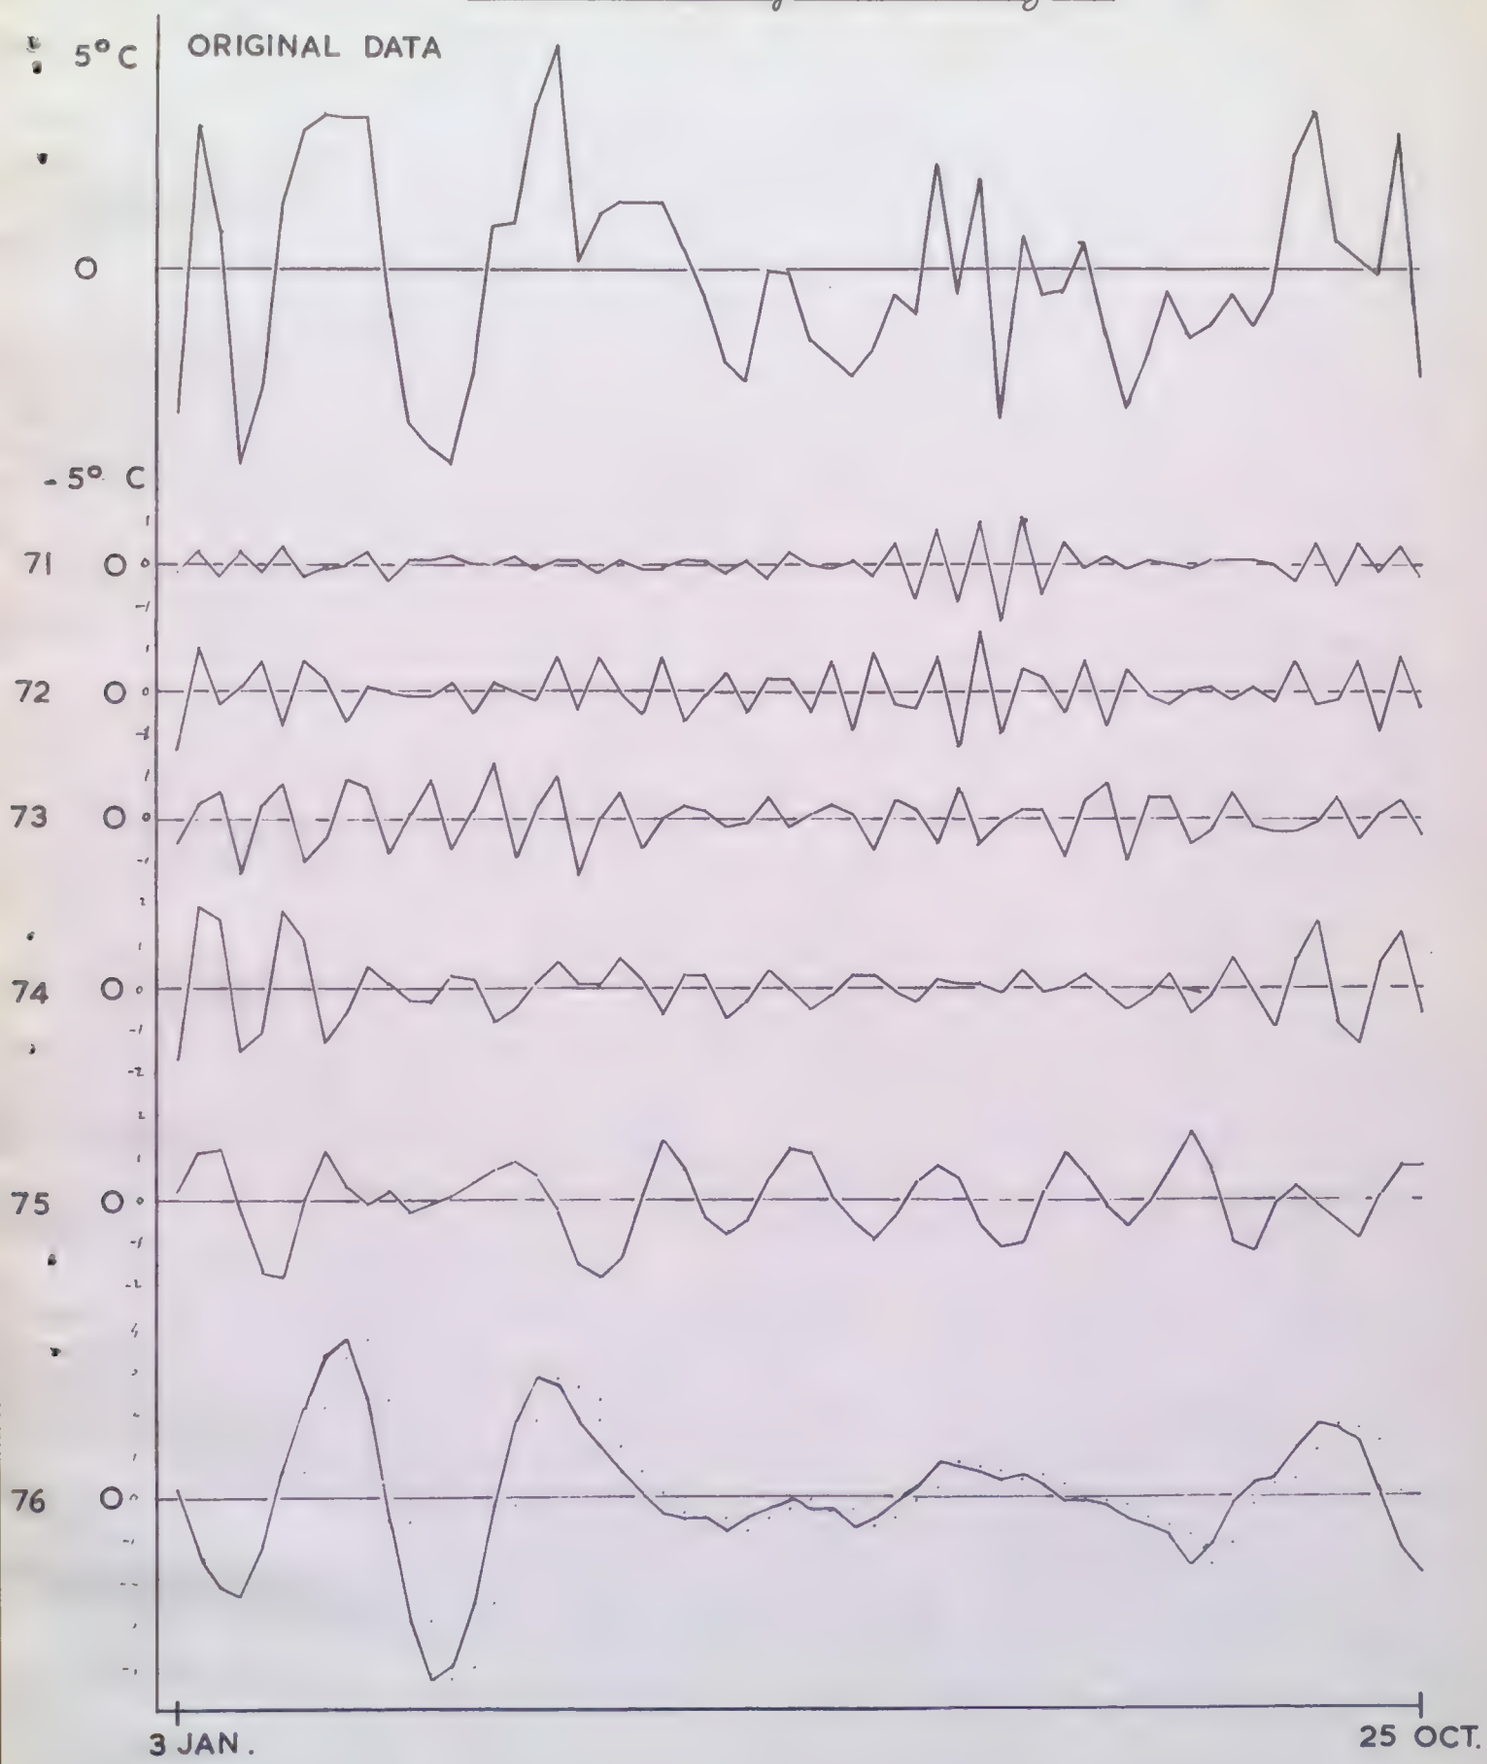

FIG. 4. SIGNIFICANCE ANALYSIS OF KEW TEMPERATURE VARIATIONS. IF A RANDOM SERIES WERE ANALYSED THE POINTS WOULD LIE ON A HORIZONTAL LINE.

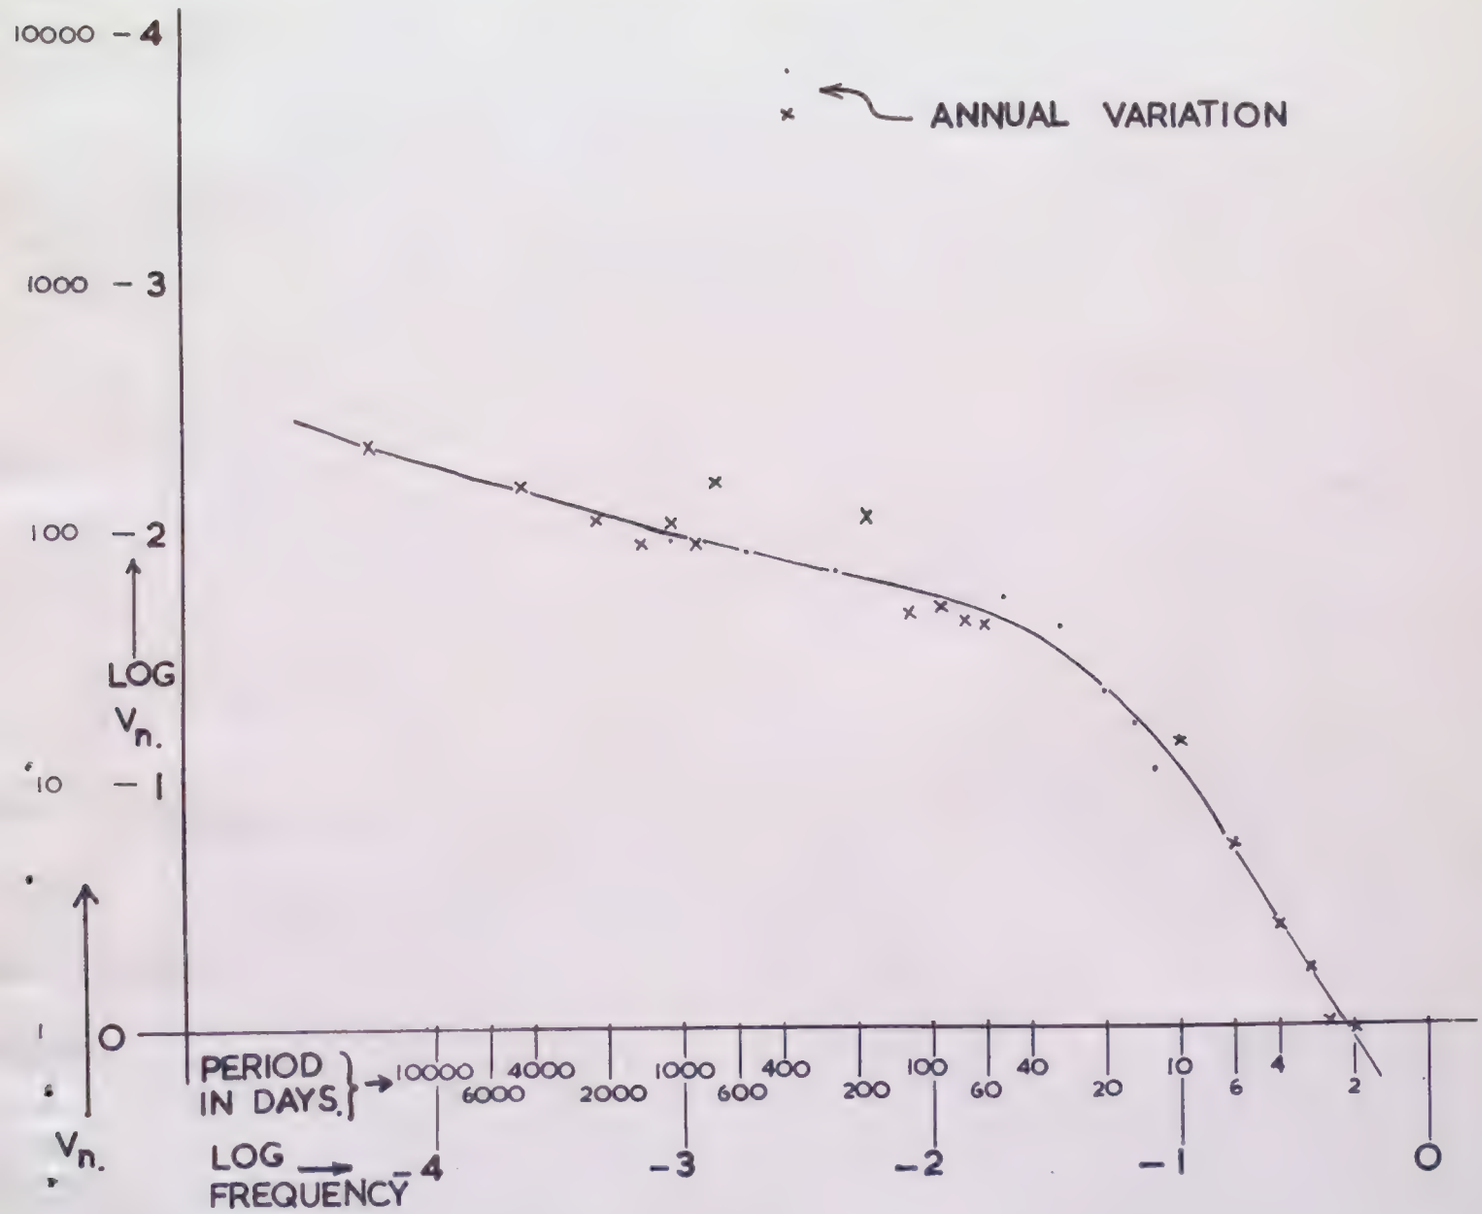

FIG. 5. POWER SPECTRUM ANALYSIS FOR KEW TEMPERATURES. ORDINATES ENCLOSING EQUAL AREAS UNDER THE GRAPH INCLUDE FREQUENCY BANDS PASSING OSCILLATIONS OF EQUAL AMPLITUDE.
The shaded area represents a variance of 1°C²

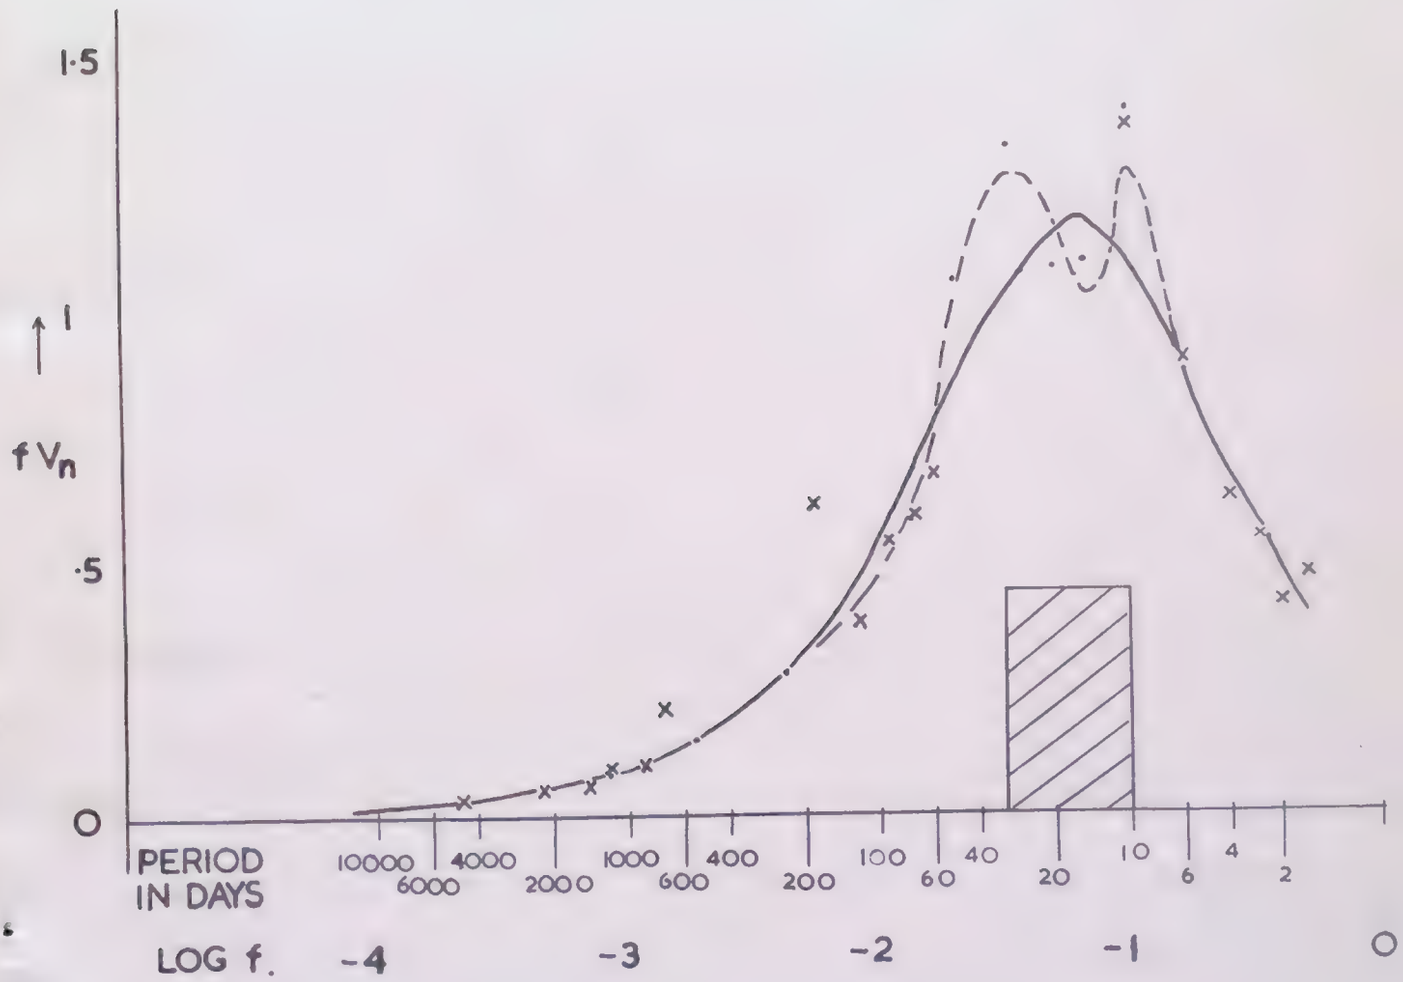

FIG. 6. KEW DAILY TEMPERATURES FOR 1955 AFTER FILTERING TO RETAIN ONLY A FREQUENCY BAND CENTRED ABOUT 1/10 (I.E. A PERIOD OF 10 DAYS)

M.R.P. 1051

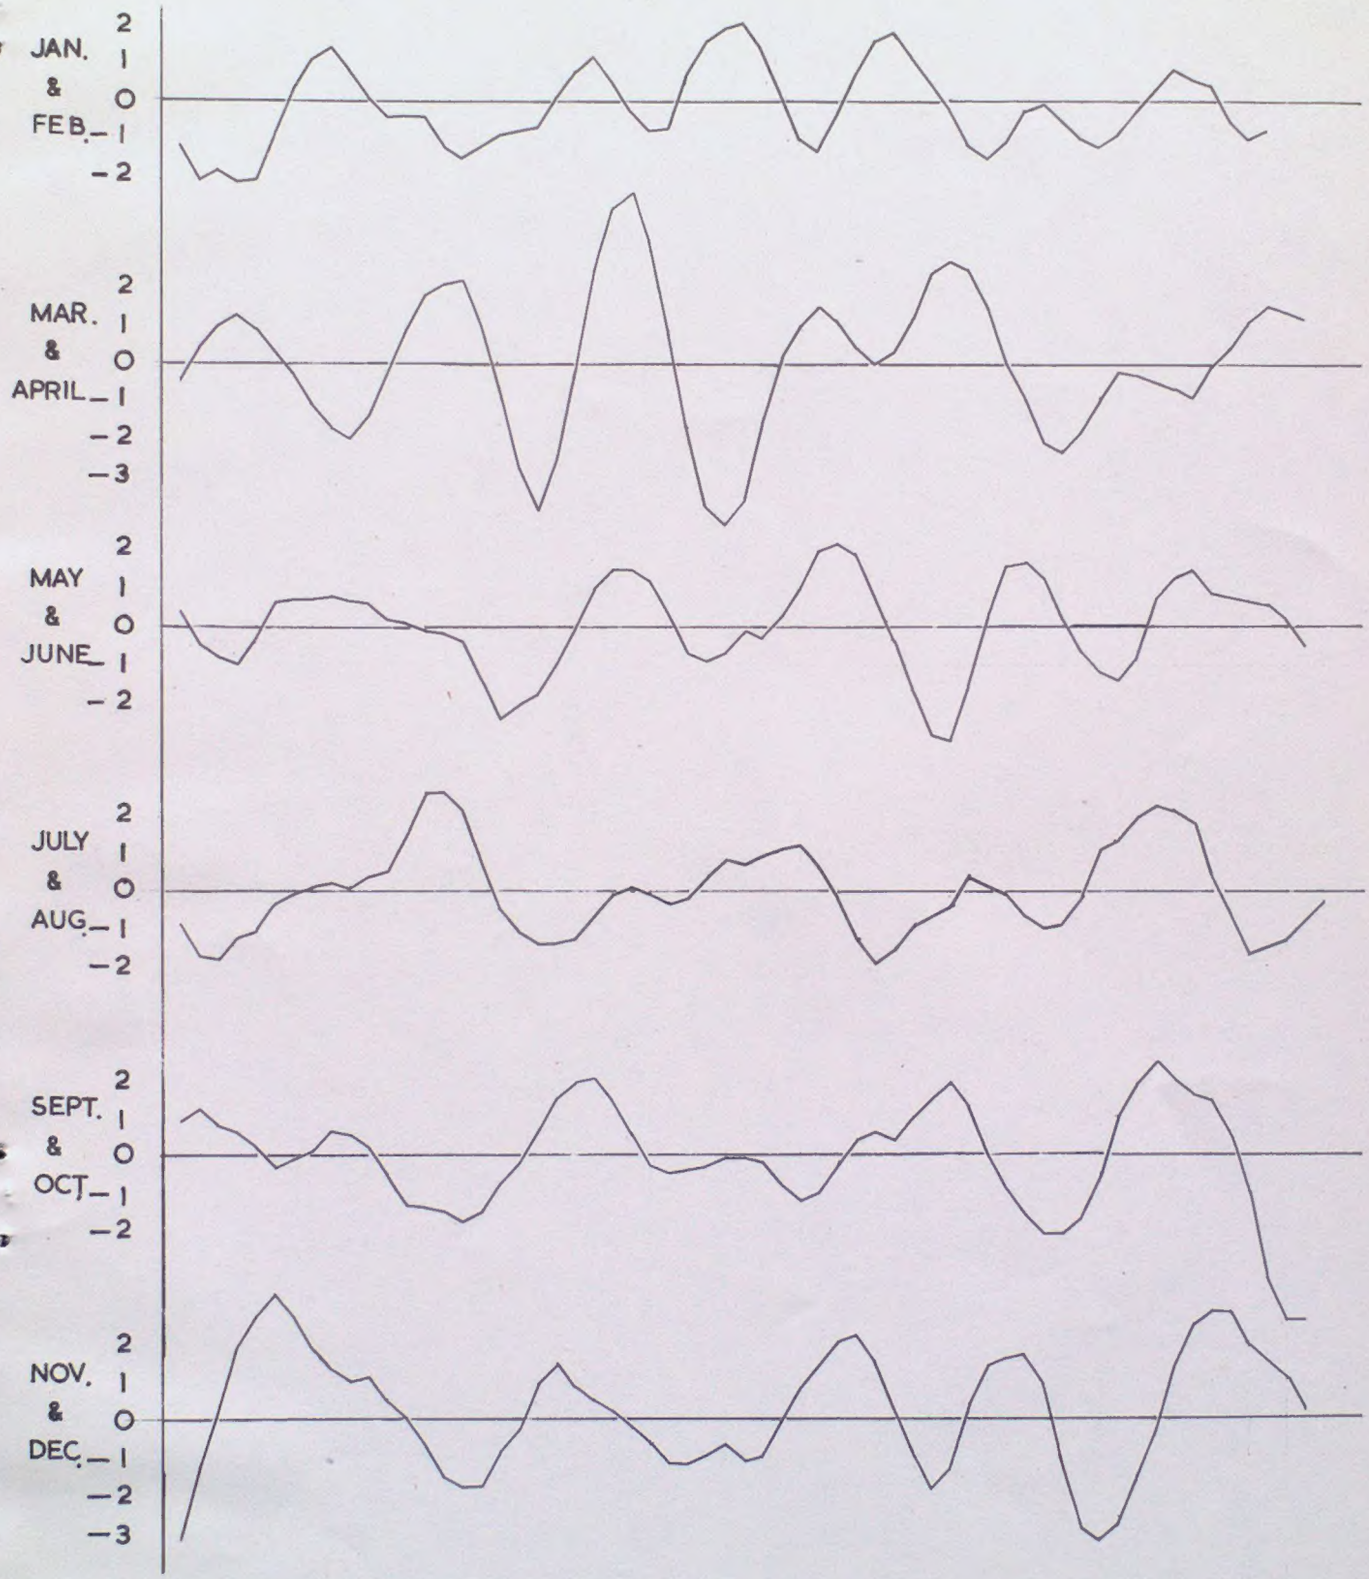

FIG. 7. KEW 5-DAY MEAN TEMPERATURES FOR 1946 - 1950 FILTERED TO SHOW THE 30-DAY WAVE IN ITS TRUE PHASE AND AMPLITUDE.

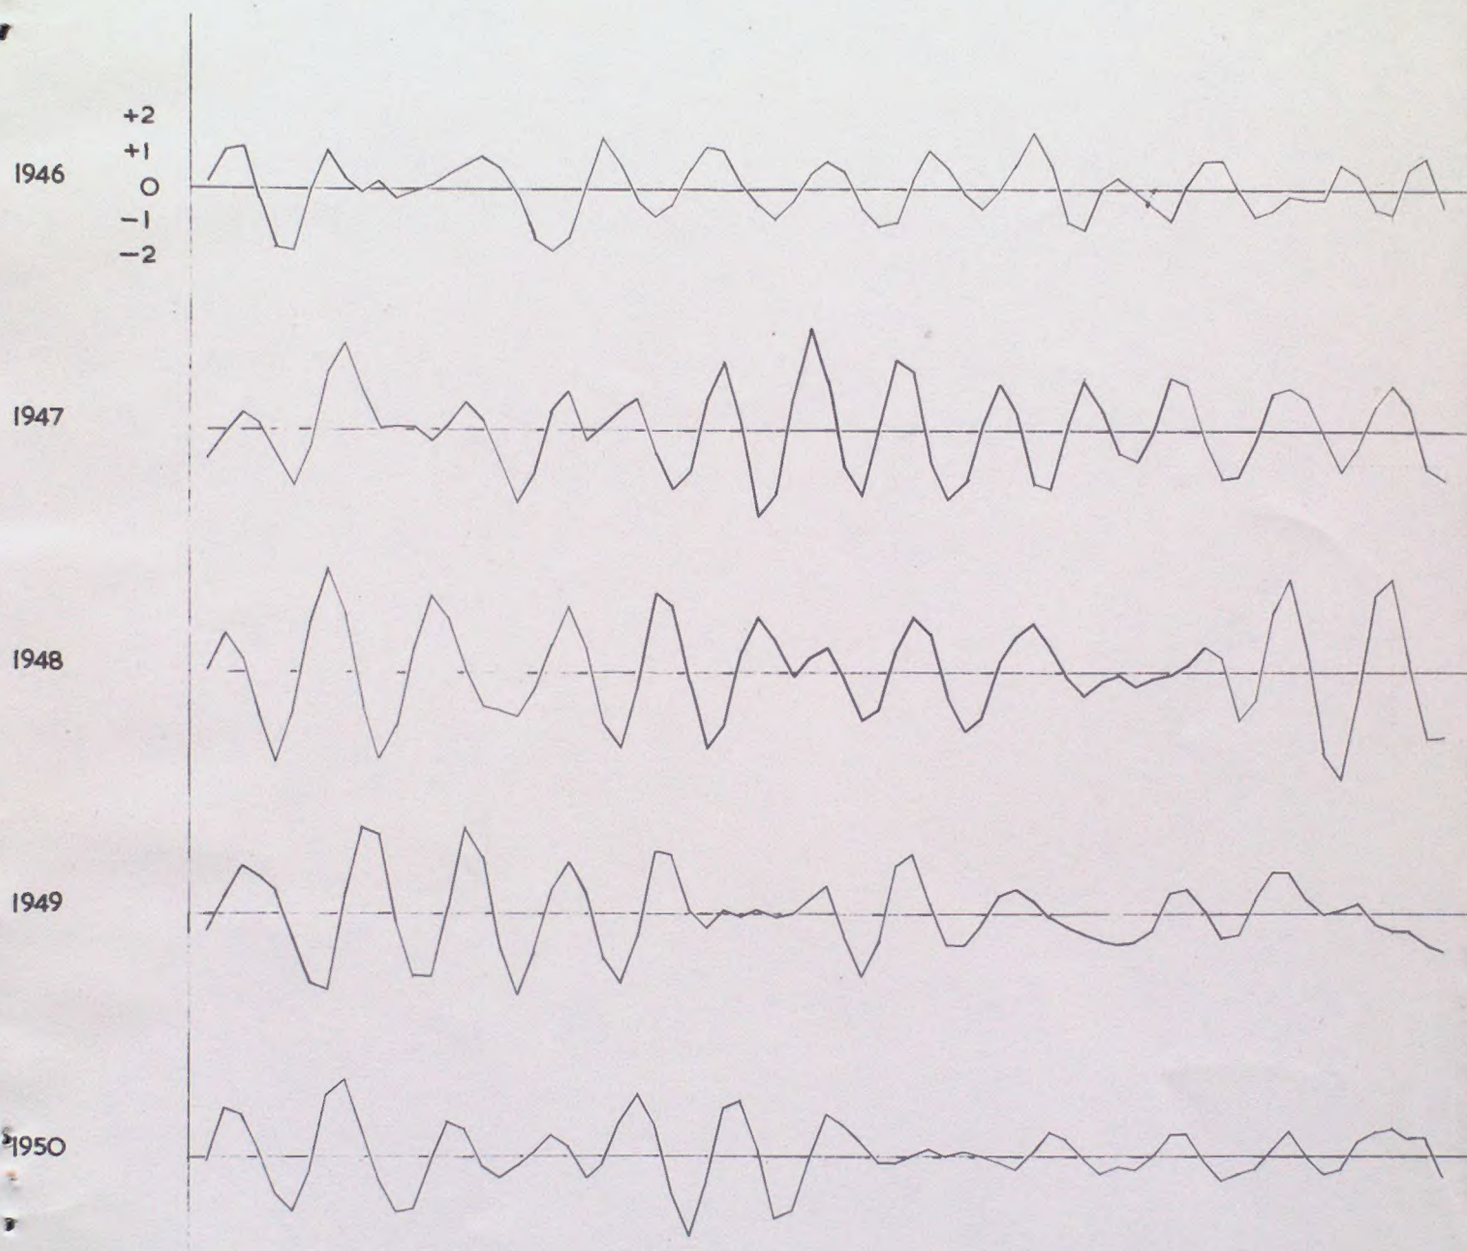

FIG. 8 BEAR ISLAND MONTHLY TEMPERATURE
FILTERED TO SHOW INTER-ANNUAL VARIATIONS

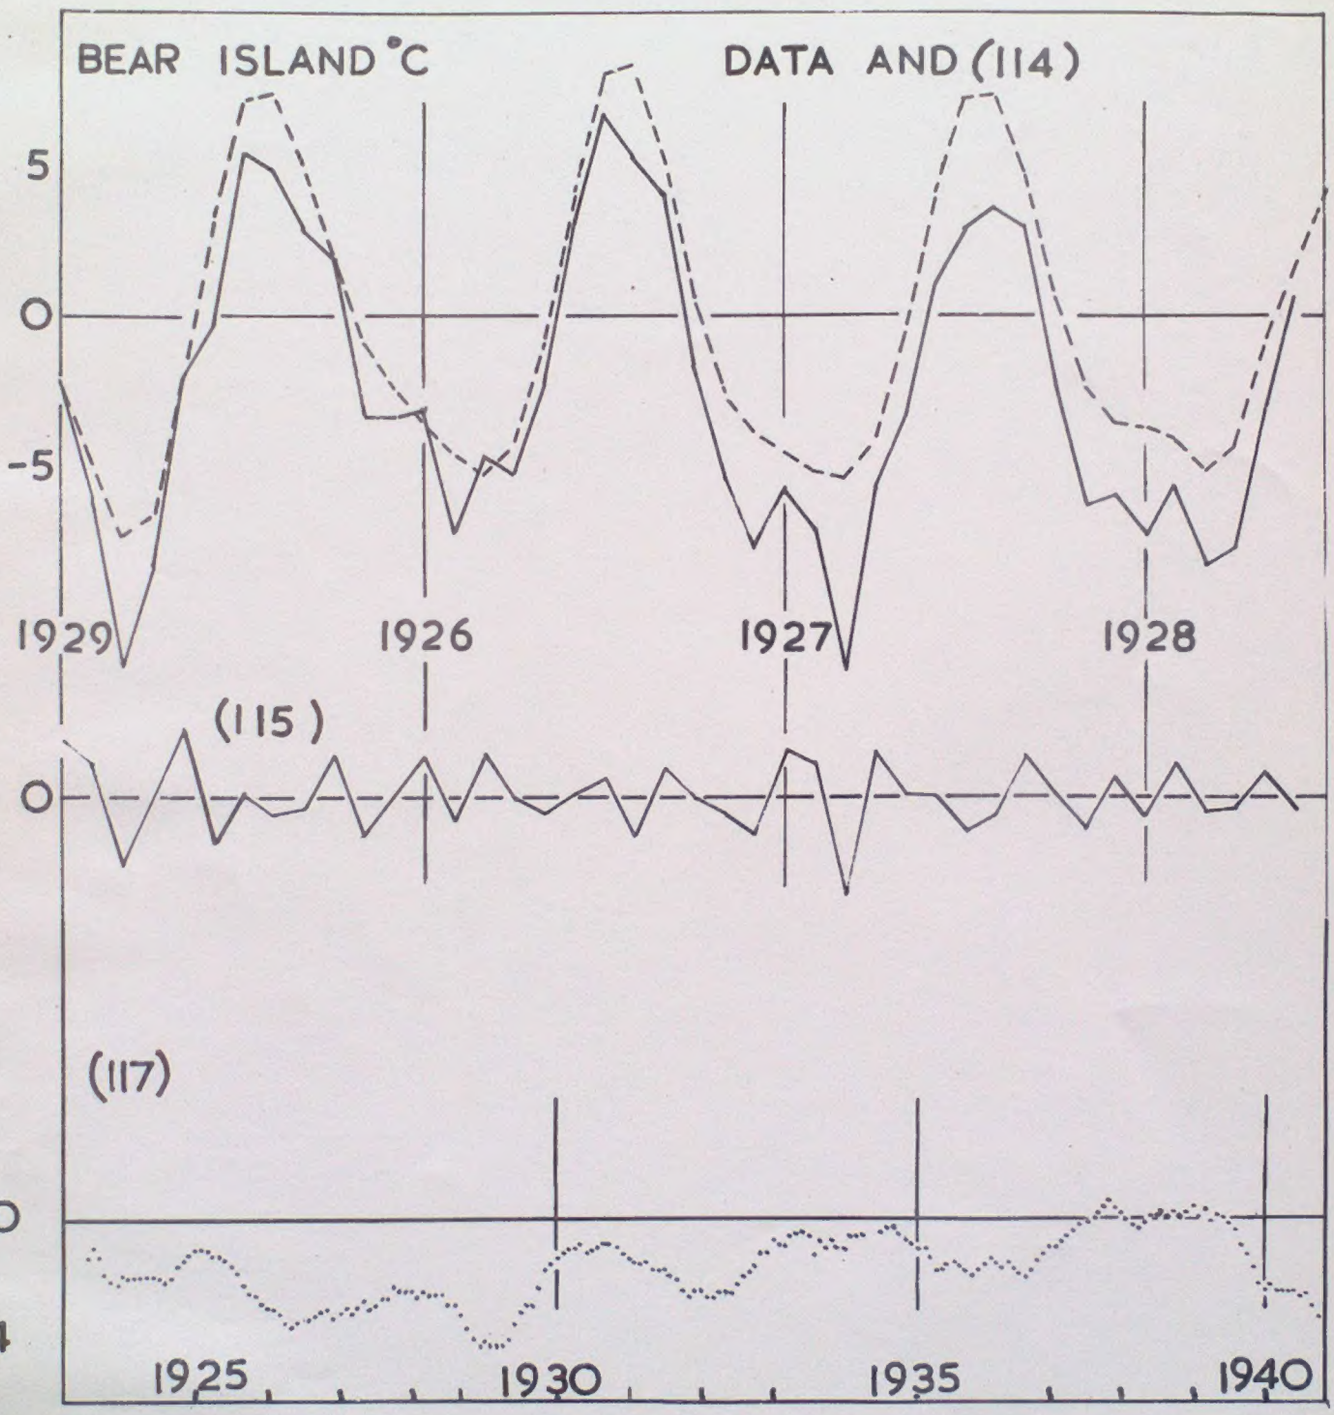

Fig. 9

M.R.P. 1051
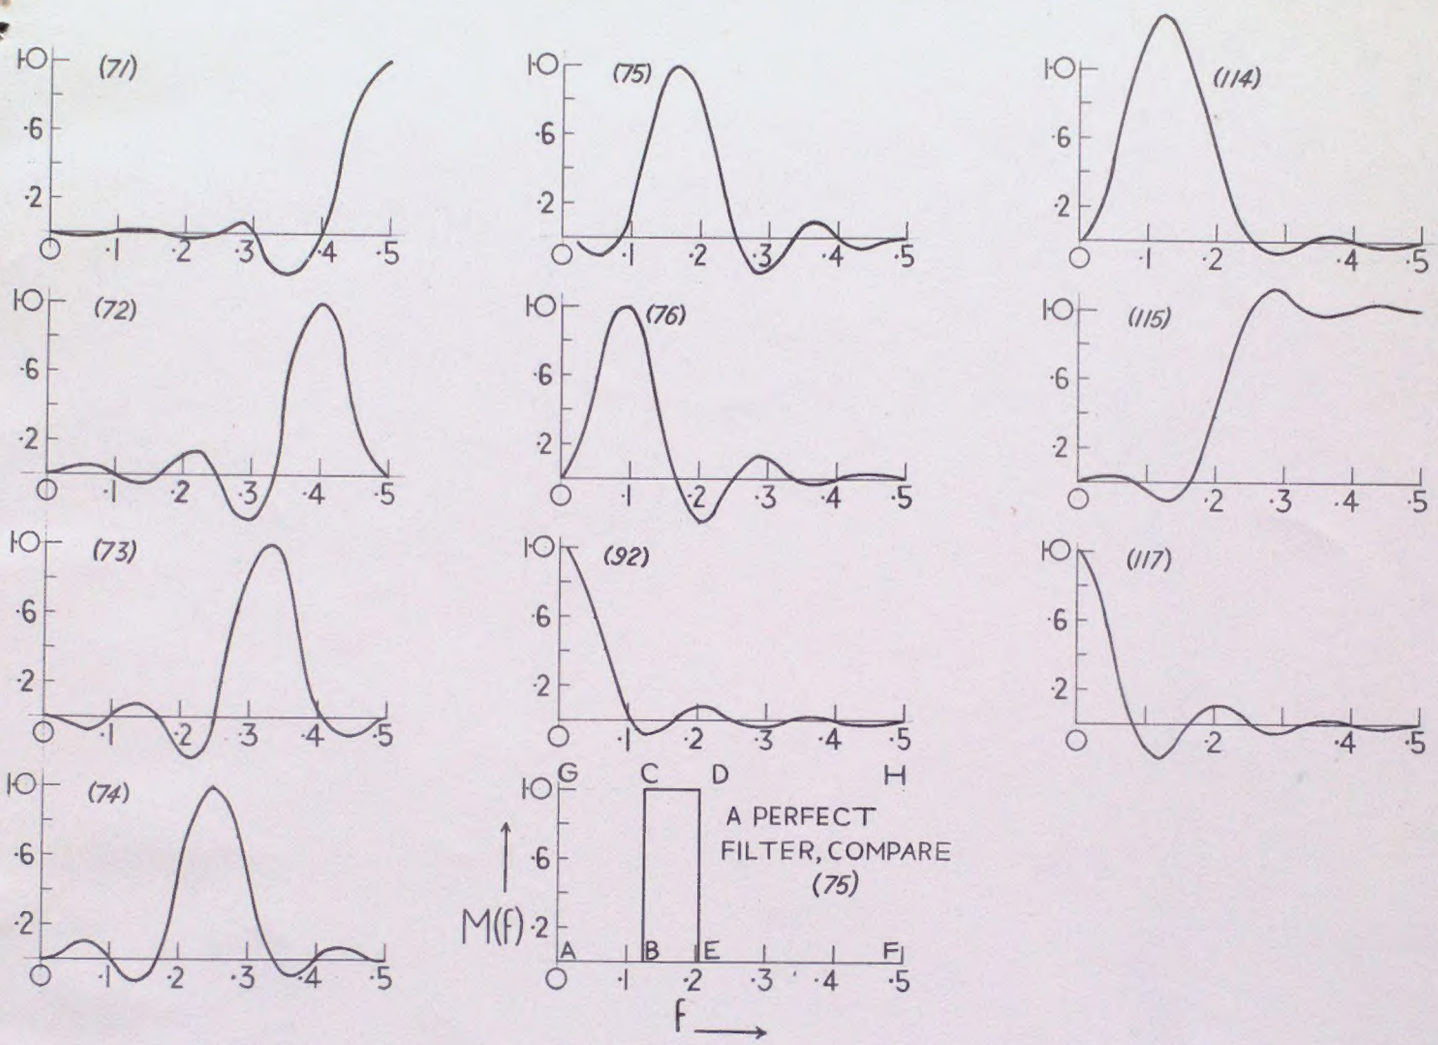
Response curves of typical filters. Filter numbers are bracketed. In all the graphs "f" is abscissa and "M(f)" ordinate